

EXPR*e*S- Real-Time VLBI

- The Current State of the Art and Next Steps

T. Charles Yun


Project Manager

JIVE, Coordinating Institution



Overview

- Project Description, context, background
- Current Status and Activities
- Possibilities and the future



Management and Overview

EXPR_eS- the Project

- EXPR_eS = Express Production Real-time e-VLBI Service
- 3 year project, start March 2006,
 - funded by FP6 DG-INFSO
 - Contract #026642.
- Objective: to create a distributed, large-scale astronomical instrument of continental and inter-continental dimensions.
- Means: high-speed communication networks operating in real-time and connecting some of the largest and most sensitive radio telescopes on the planet.

EXPR_eS' Goal

The overall objective of EXPR_eS is to create a **production-level, real-time**, “electronic” VLBI (e-VLBI) **service**, in which the radio telescopes are reliably **connected** to the central **supercomputer** at JIVE in the Netherlands, via a high-speed optical-fibre communication **network**...

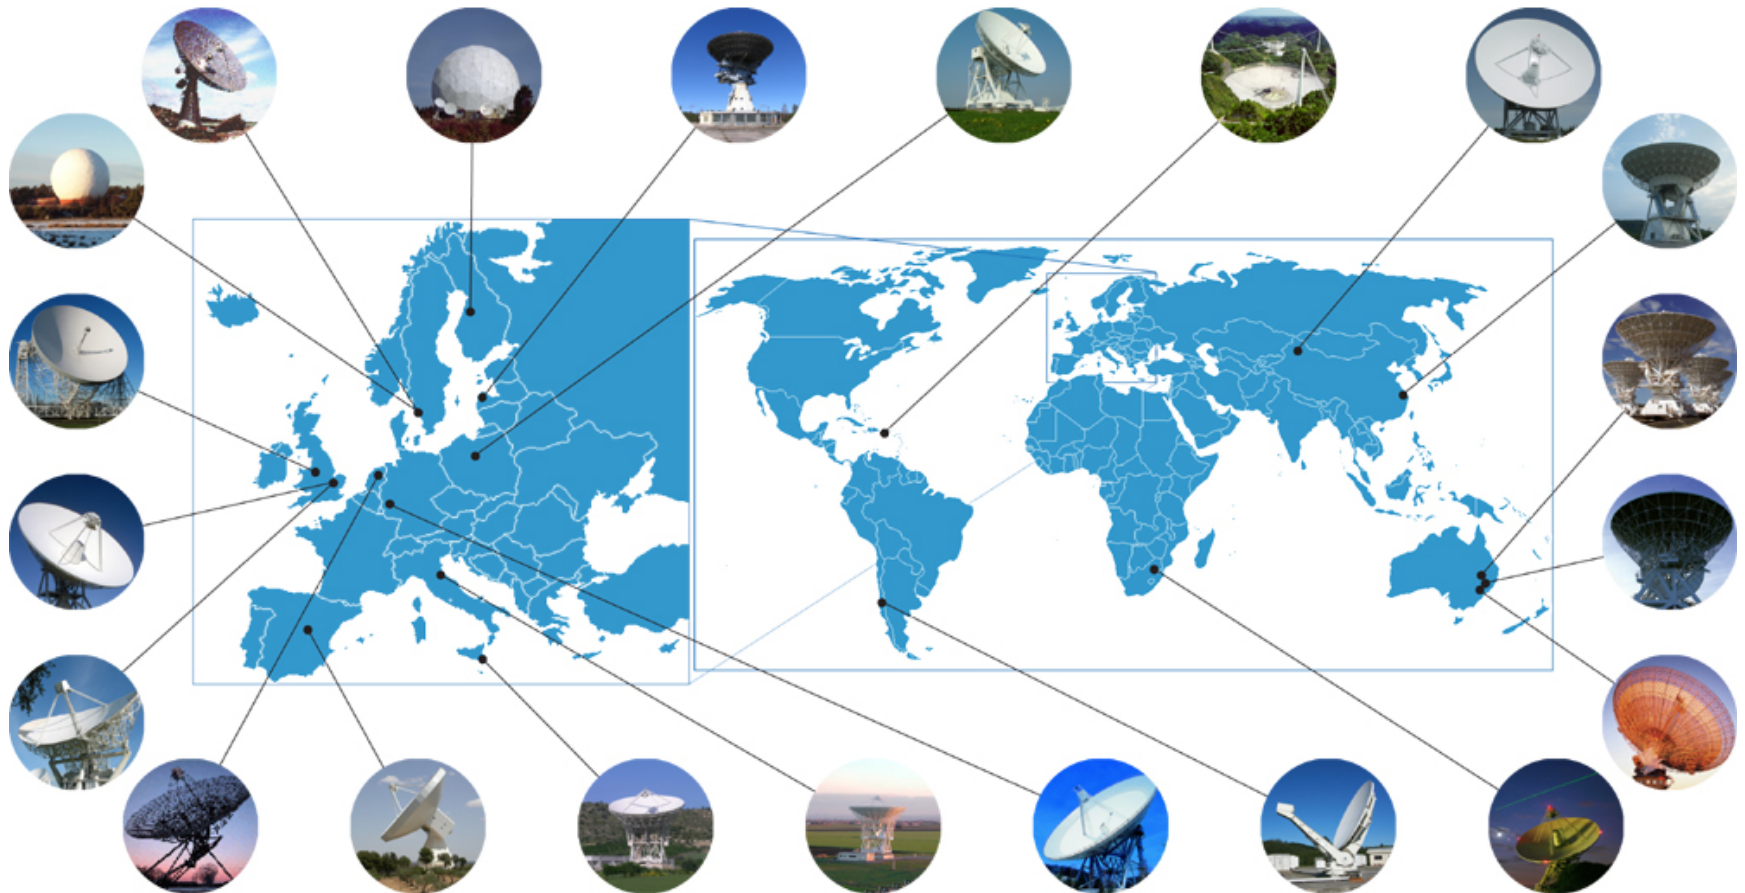
- or -

Make e-VLBI *routine, reliable* and *realistic* for astronomers.

EXPR_eS Partners

- Joint Institute for VLBI in Europe (coordinator)
- AARNET Pty Ltd., Australia
- ASTRON, the Netherlands
- Centro Nacional de Informacion Geografica, Spain
- Chalmers Tekniska Hoegskola Aktiebolag, Sweden
- Commonwealth Scientific and Industrial Research Organization (CSIRO), Australia
- Cornell University, USA
- Delivery of Advanced Network Technology to Europe Ltd. (DANTE), UK
- Istituto Nazionale di Astrofisica, Italy
- Instytut Chemii Bioorganicznej PAN, Poland
- Max Planck Gesellschaft zur Foerderung der Wissenschaften e.V., Germany
- National Research Foundation, South Africa
- Shanghai Astronomical Observatory, Chinese Academy of Sciences, China
- SURFNet b.v., The Netherlands
- Teknillinen Korkeakoulu, Finland
- The University of Manchester, UK
- Universidad de Concepcion, Chile
- Uniwersytet Mikołaja Kopernika, Poland
- Ventspils Augstskola, Latvia


Telescope Locations



Copyright: EXPReS
Telescope photos used with permission.

Radio Astronomy, VLBI

- Single radio-telescopes severely limited in resolution ($\propto \lambda/D$)
- Can be overcome by interferometry, creating a virtual telescope
- In VLBI the telescopes can be distributed on continent scales
- Sensitivity depends on bandwidth sampled and number of bits transported
 - And the telescope diameter



Jodrell Bank Lovell (UK)



Onsala 20m (SE)



Onsala 25m (SE)



Dwingeloo (NL)



JOINT INSTITUTE FOR VLBI IN EUROPE

Jodrell Bank Mk2 (UK)



Metsähovi (FI)



Torun (PL)



Urumqi (CN)



Arecibo (PR)



Cambridge (UK)



Wettzell (DE)



Seshan (CN)



Hartebeesthoek (ZA)



Westerbork (NL)



Effelsberg (DE)



Medicina (IT)



Robledo (ES)



Yebes (ES)




Noto (IT)



The “problem” with traditional VLBI

- Telescopes collected data on tapes... heavy and bulky... postal mail... once all the tapes arrived... tapes were lost/damaged... hard drive arrays slightly improved the situation...
- e-VLBI solved shipping problem... introduced regular and flexible upgrades to data transport process
- It will soon (months) be possible to send data faster than the correlator can process



Why e-VLBI?

- e-VLBI:
 - shortens the delay between observations and images
 - From months to hour
 - Can offer more frequent and reliable observations
 - Potentially more bandwidth \times observation time
 - Ensures data quality already during observations
 - Relaxes storage logistics for raw noise data
 - Enables a new class of fast response observations
- Intellectual collaboration across disciplines
 - Pushing the technology, protocols etc
 - International cooperation, infrastructure investments


Operations

We run a PRODUCTION Service


- Monthly observations
- Published science from our observations
- e-VLBI is competitive with disk based VLBI
 - network bandwidth is on par with and will soon overtake (current) disk recording speeds
 - Stability of e-VLBI correlation is greater than that of disk based correlation
 - Longest uninterrupted correlation is more than 12 hours
 - Stopped only to look at new source
 - Large dishes now online- sensitivity
- All the benefits of traditional VLBI, but now with immediate monitoring and feedback

512Mbps Production

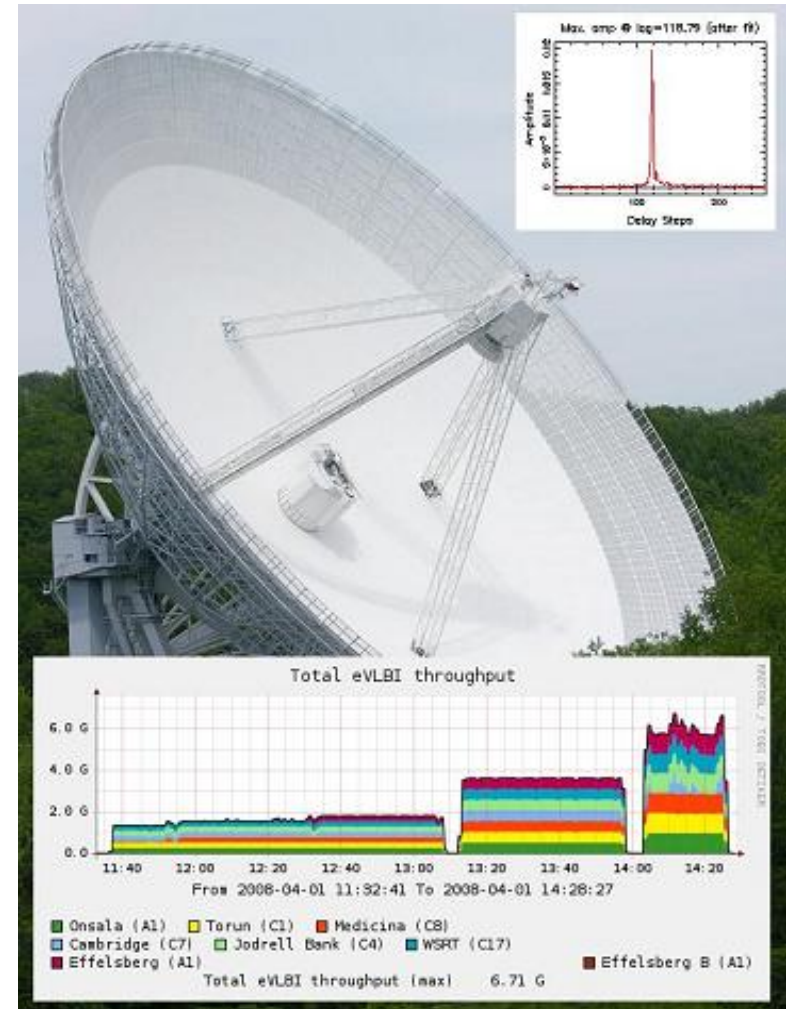
- On 8-9 April the first production 512 Mbps e-EVN observations took place
- Two science projects were observed
 - The first was a normal non-triggered project on Arp229 (RP009)
 - The second (RT006A) was triggered on Cyg X-3
- The correlation of RP009 started at 14:37 UT and ended at 3:30 UT
 - Longest ever uninterrupted correlation job (**including disk operation**)



Recent results




- First connections to TIGO, Chile
- Much improved connection to Arecibo, Puerto Rico



Effelsberg on line (at 1 Gbps!)

TNC 2008 Demo

- Demonstrated simultaneous operations on 4 continents
 - Arecibo (Puerto Rico)
 - Effelsberg (Germany), Westerbork (Netherlands), Onsala (Sweden), Medicina (Italy)
 - Hartebeesthoek (South Africa)
 - TIGO (Chile)
- Live demo on stage at TERENA conference



Yebes, new 40m telescope

- First fringes last month
- first ever VLBI participation
- In so-called ftp tests
- Using software correlator
- First ever telescope to have first fringes without recording?



Network Infrastructure

EXPReS Telescope Locations

The *EXPReS* network



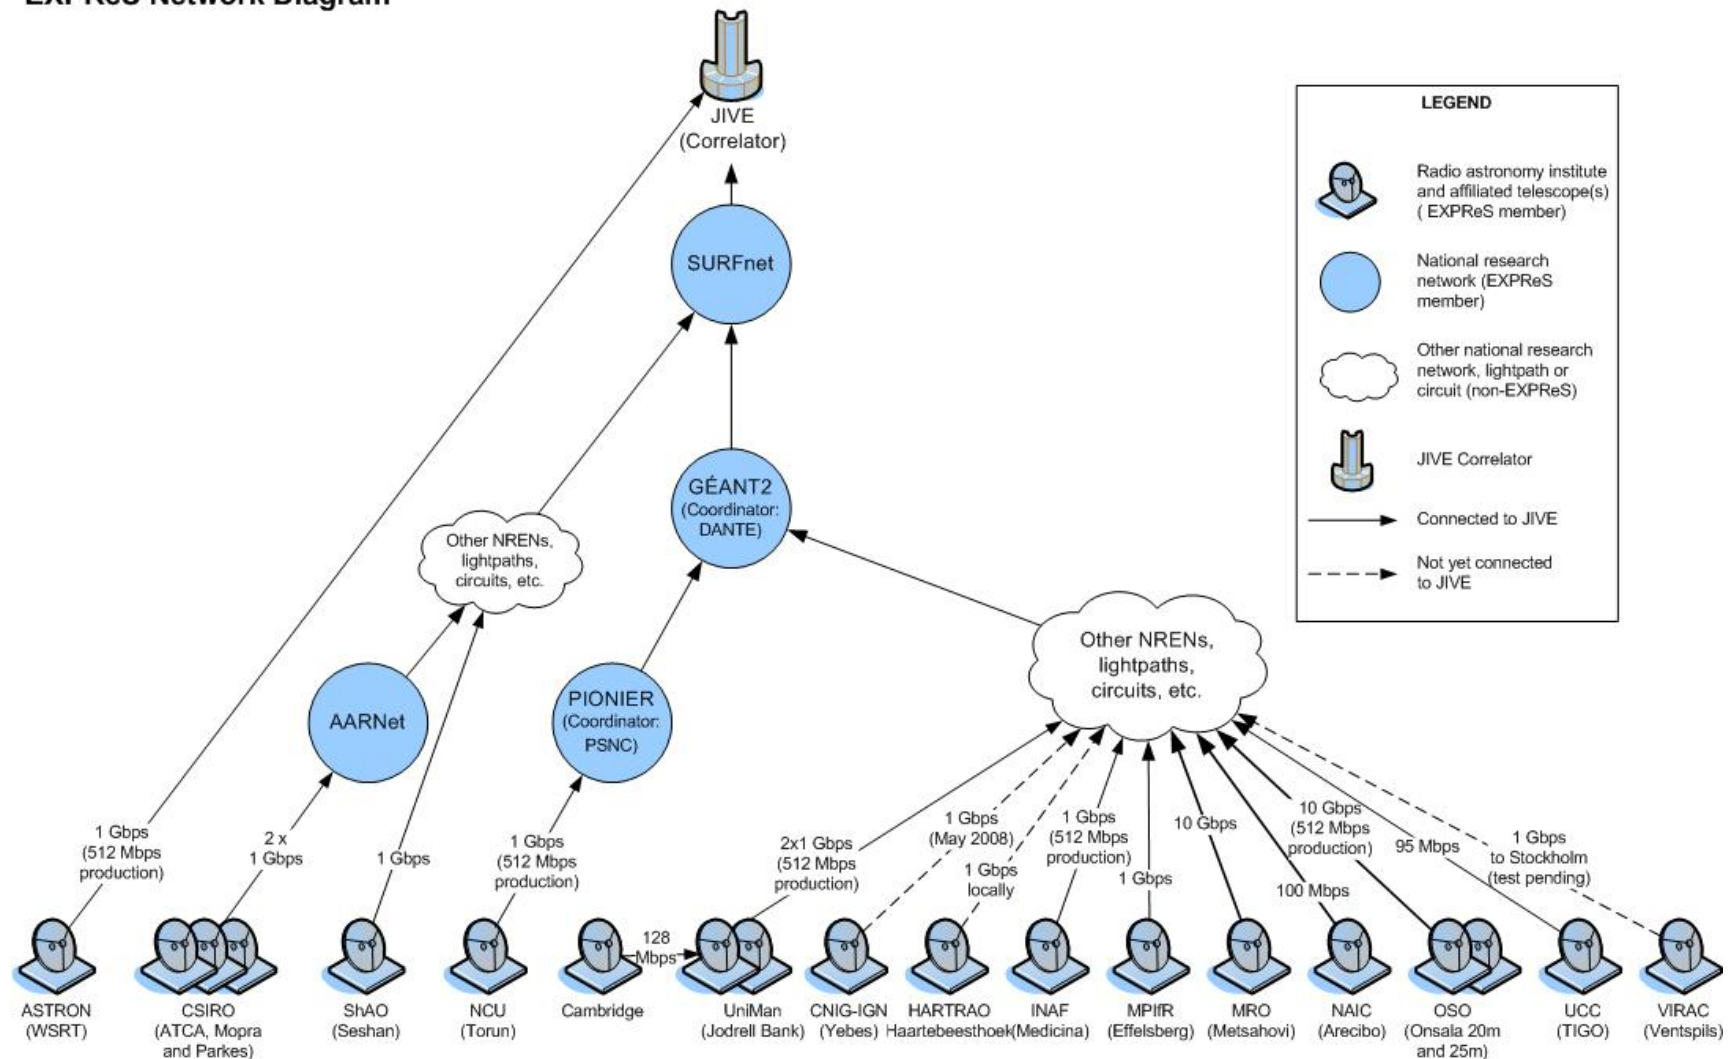
Network status as per 2008-05-02. Image created by Paul Boven <boven@jive.nl>. Satellite image: Blue Marble Next Generation, courtesy of Nasa Visible Earth (visibleearth.nasa.gov).

Quick Summary of EC Contribution

- Funding institutions to {connect|upgrade} network connections for telescopes: last mile, new hardware, lease costs, etc.
- Each institution is responsible for spending and justifying expenses
- First year was a learning experience, in particular for non-EU partners (China, Australia, etc)
- EC partial funding (8.75:1 in proposal) acts as a great catalyst of work and national funding
- Very uneven cost in Years 1-3 across partner institutions

Network Connections to JIVE

EXPRéS Network Diagram




Telescope Connection Speeds

Telescope	Current BW	Expected BW	Year	Notes
JIVE correlator	8 x 1 Gbps + 1 x 5 Gbps	16 x 1 Gbps	2008	connected
WSRT (14x25m)	1 Gbps			connected
Onsala (20+25m)	10 Gbps			connected
Jodrell Bank (76m)	2 x 1 Gbps	10 Gbps	2008	connected
Cambridge (32m)	1 Gbps			connected
Torun (32m)	1 Gbps			connected
Metsähovi (14m)	10 Gbps			connected
CNIG-Yebes (40m)	2 Mbps	1 Gbps	2008	
Effelsberg (100m)	1 Gbps	4 Gbps		connected
Medicina (32m)	1 Gbps			connected
Noto (32m)			unknown	
Sardinia (64m)		2,5 / 10 Gbps	2009	
Shanghai (25m)	500 Mbps	1 Gbps	2008	connected
Urumqi (25m)		1 Gbps	2009	
Miyun (50m)	35 Mbps	Not possible		No longer considered
Yunnan (10m)	35 Mbps	Not possible		No longer considered
VIRAC (32m)	1 Gbps			connected (test pending)
Hartebeesthoek (26m)		1 Gbps	unknown	connected locally
Tigo (6m)	95 Mbps	155 Mbps	2008	512 Mbps (for tests)
Arecibo (305m)	100 Mbps	1 Gbps	2008	512 Mbps (for tests)
ATNF/CSIRO (Parkes, Mopra, ATCA)	2 x 1 Gbps			connected

note: already out of date

Correlator R&D


Correlation



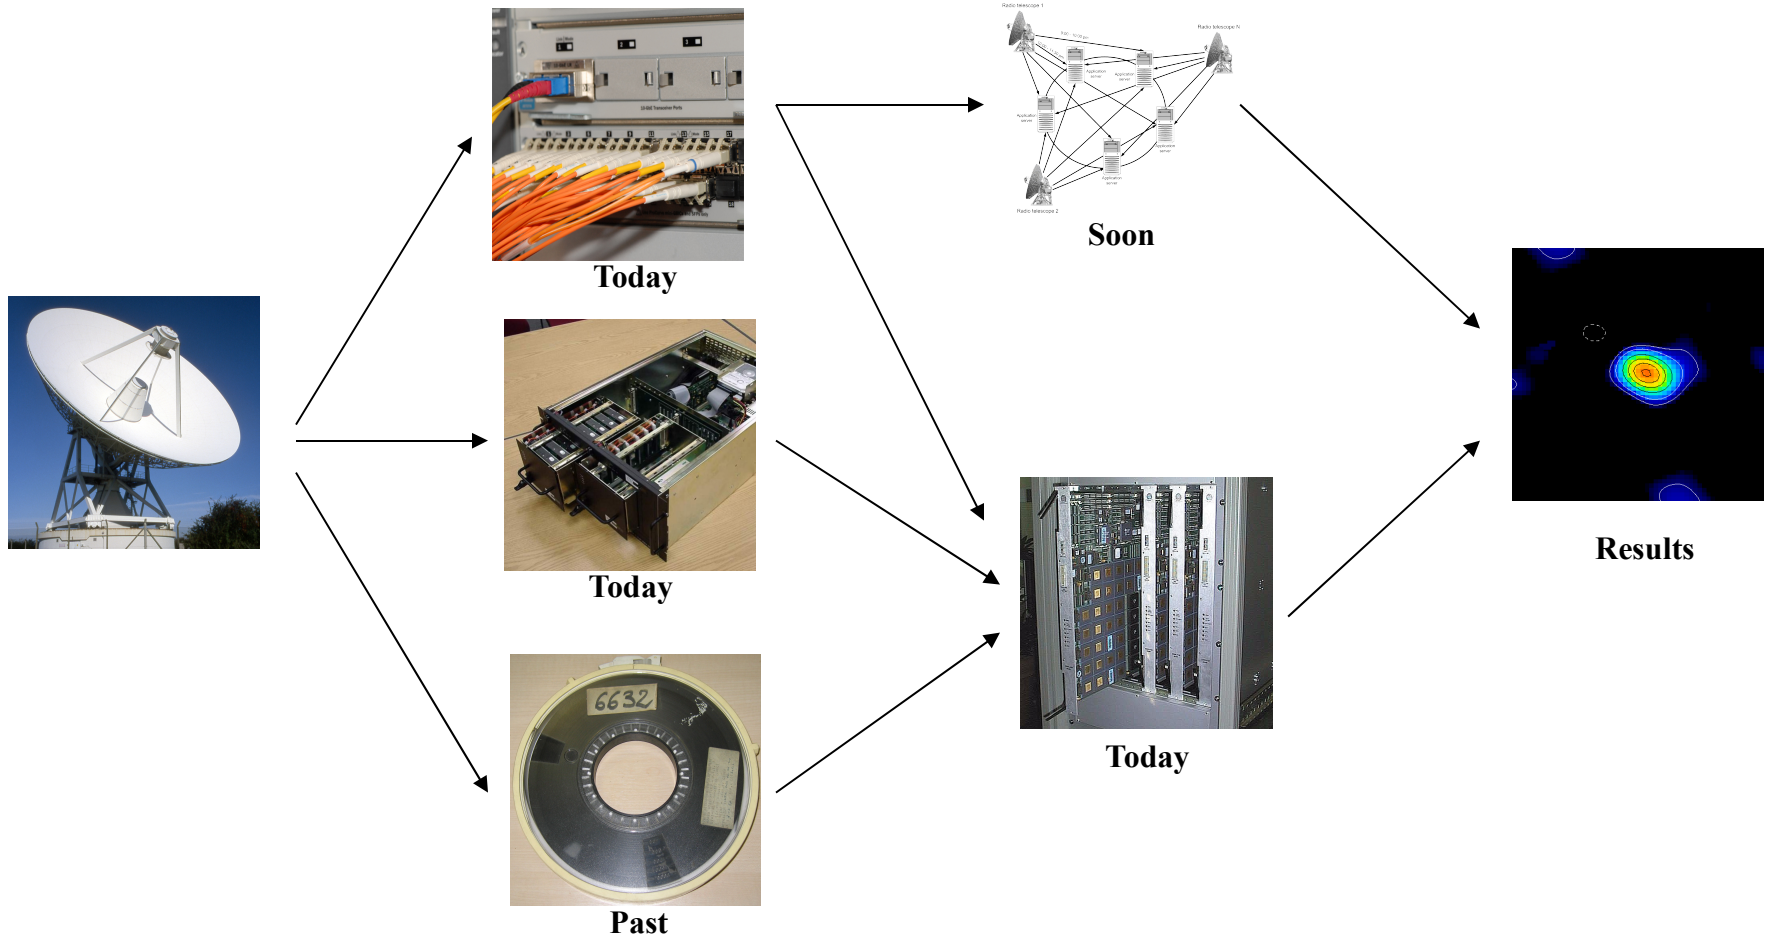
- Data between all telescope pairs needs to be correlated
 - Requires accurate timing to line up the data signals
 - Dedicated supercomputer aligns signals, evaluate correlation function, accumulate data
- Central correlator for EVN is at JIVE, Dwingeloo
 - Based on 90's VSI technology, deals with 1 Gbps data

Why Distributed/Software Correlation?

- Cost to build correlator...
limited flexibility (majority of computation in custom hardware)... preset data input rates... scheduling of scarce resource (correlator)... upgrade cost forces longer life-cycle than desired



Basic Overview




Existing Correlators

Location	# Stations	Speed	Type
JIVE		16 1024 Mbps	hardware

Data distribution over grid sites (1)

Baseline slicing




Pros

- Small nodes
- Simple implementation at node

Cons

- Multiplication of large data rates, especially when number of baselines is large
- Data logistics complex
- Scalability complex

Data distribution over grid sites (2)



All data to one site


Pros

- Simple data logistics
- Central processing
- Live processing easy
- Slicing at the grid site
- Dealing with only one site.

Cons

- Powerful central processing site required

1



All data to different sites


Pros

- Smaller nodes
- Live processing possible
- Data slicing at nodes

Cons

- Multiplication of large data rates
- Simultaneous availability of sites when processing live

2



Time slicing


Pros

- Smaller nodes
- Smaller data rates
- Simple implementation
- Easy scalable
- No data multiplication

Cons

- Complex data logistics after correlation
- Live correlation complex

3



Channel slicing

Pros

- Smaller nodes
- Live processing per channel
- Simple implementation
- Easy scalable

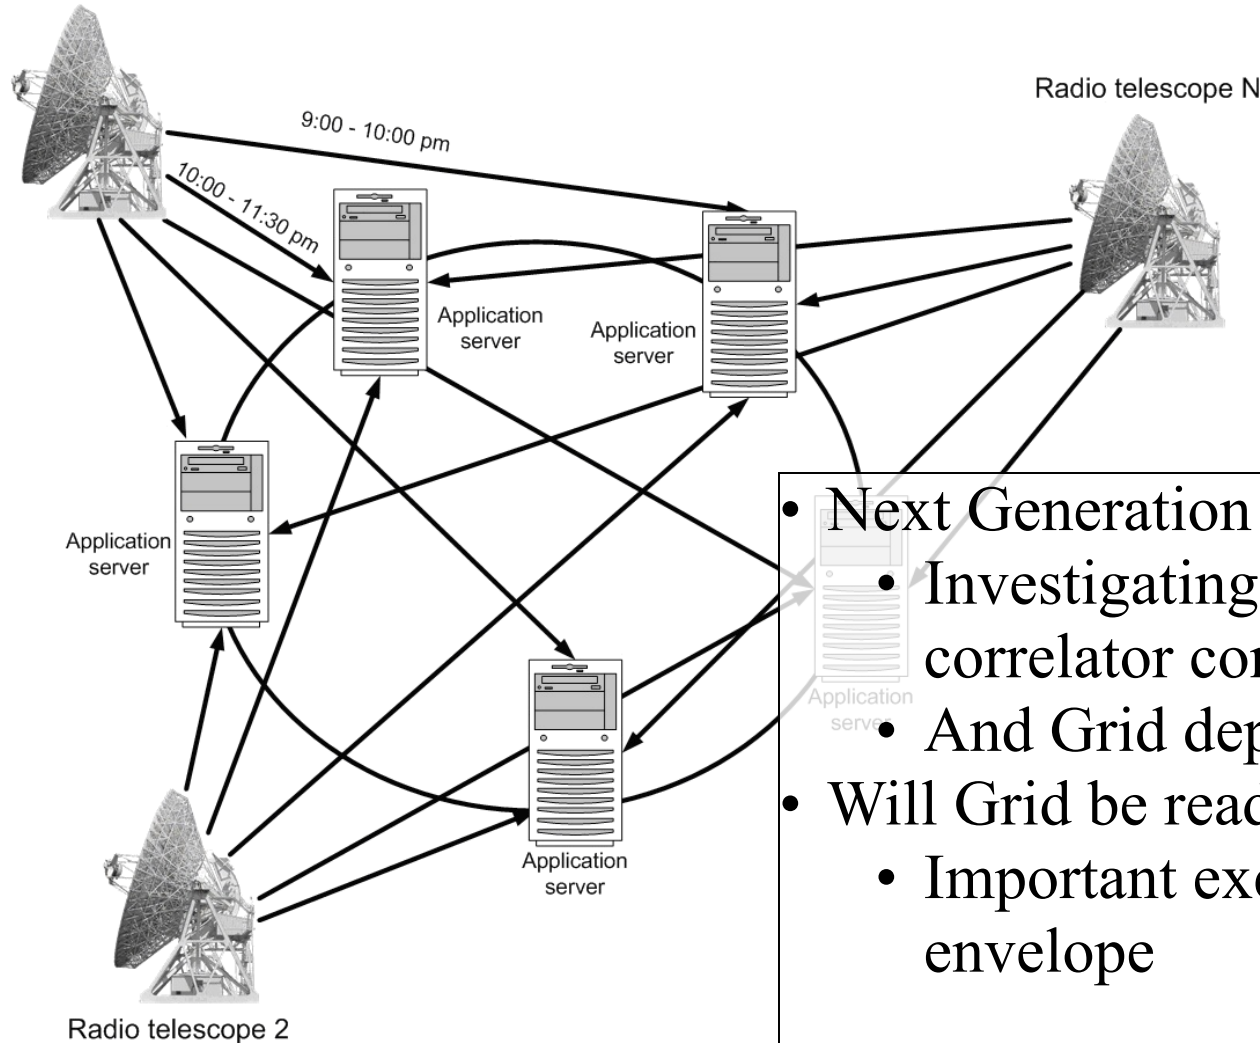
Cons

- Channel extraction at telescope increases data rate

4

Distributed Correlation

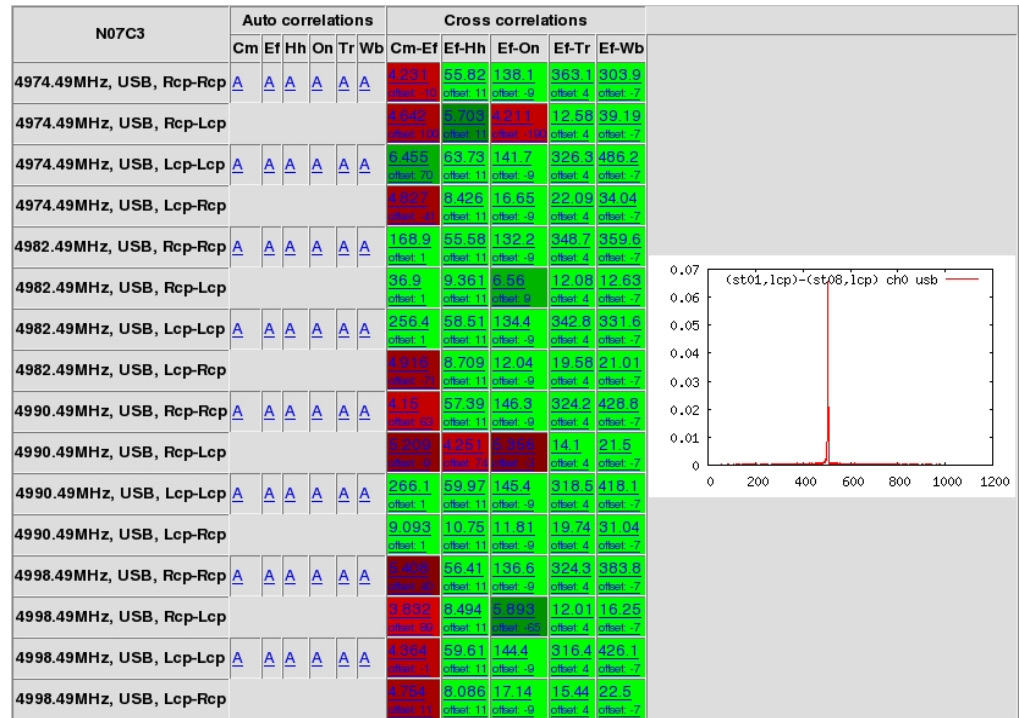
Radio telescope 1



- Next Generation Correlator
 - Investigating distributed correlator concept
 - And Grid deployment
- Will Grid be ready for this?
 - Important exercise to push the envelope


Proving software correlator

- FTP Fringe tests: short slices of an experiment to test equipment.
- Cluster running NICT correlator broke down in May 2007
- SFXC was successfully used for the last three FTP Fringe Tests





Correlator Results

- Direct comparison of software correlator output with the EVN data processor at JIVE (hardware) for data from a K-band NME



Benchmarks on DAS-3

- 40% of realtime (4 telescopes, 256Mbps each) on 30 quad-cores
- bottleneck is CPU power
- GPU's/FPGA's could enhance performance in same topology
- Using the algorithm originally aimed at Spacecraft navigation
- Targeting AutoBAHN



Closing thoughts

Working with the EC

- EC Project Accounting Rules
 - Complicated- and no one has good answers
 - Requires a great deal of paperwork
 - Audits, accountants, european style finances
 - Projects should have a local “expert” on EC rules
 - An additional contact from the coordinating institute
 - Direct access to the EC to obtain answers (generally a finance support person, not the science officer)
- Project reports and deliverables
 - EXPReS is required to submit monthly reports to the EC with updates from each activity
 - Annual review documents require considerable effort from everyone

Current Challenges

- Correlation is **extremely** CPU intensive
 - Optimizations/shortcuts in algorithm?
 - Bigger clusters, Clusters of clusters
 - Moore's Law (in 18 months, the same code is twice as fast)
- Hybridization
 - FPGAs, FPGA + Cluster,
 - Inexpensive PS3/Cell processor cluster (probably unrealistic, but interesting potential)
- Cost of networking
 - Local Loop/Last mile, regional
 - International
 - Underwater/Long-haul
- Most institutes lack breadth of skill outside of astronomy
 - Network, Computational, IT/sysadmin, software

Next Steps

- Most Important Point
 - no money right now
- Looking for interesting ideas
 - ideas -> partners -> projects -> money
- Collaboration will span multiple disciplines
 - Astronomy, computer science, international networking, hardware design
- Hoping that you are interested and will contribute/participate in the efforts

Questions/Answers

- Contact information
T. Charles Yun
Project Manager
EXPR_eS (JIVE)
tcyun \at\ jive dot nl
- Additional Information
<http://expres-eu.org/>
<http://www.jive.nl/> [note: only one “s”]
- EXPR_eS is made possible through the support of the European Commission (DG-INF_{SO}), Sixth Framework Programme, Contract #026642



INAF




ISTITUTO NAZIONALE DI ASTROFISICA
NATIONAL INSTITUTE FOR ASTROPHYSICS



Max-Planck-Institut
für Radioastronomie



HELSINKI UNIVERSITY OF TECHNOLOGY
Metsähovi Radio Observatory



The University of Manchester
Jodrell Bank
Observatory



JOINT INSTITUTE FOR VLBI IN EUROPE

